

POWER CONTROL SYSTEM FOR SYNCHRONOUS MEMORY DEVICE

This application claims priority to the provisional patent application entitled "Power Control System for Synchronous Memory Device", Serial Number 60/061,664, filed October 10, 1997.

RELATED APPLICATION

A related application was filed on May 7, 1996, Serial No. 08/648,300, entitled "Asynchronous Request/Synchronous Data Dynamic Random Access Memory", assigned to the same assignee as this application, hereby incorporated by reference as background information.

BACKGROUND OF THE INVENTION

The present invention relates to power domains, and in particular to clock power domains in memory systems such as dynamic random access memories (DRAMs).

Power consumption is a constraint on computer systems both by virtue of limited power available in portable, battery-operated systems, and also limited heat dissipation for high power devices. As devices are made faster by increasing their clock speed, the power requirements also increase since clock signal lines, receivers, and other clock circuits consume more power and generate more heat as device clock rates increase.

Some memory systems operate asynchronously. Other memory systems, to increase the speed and bandwidth, operate synchronously using a clock signal. For these synchronous systems, clock power becomes an important issue at high frequencies. High power consumption by the clock signal can exceed thermal cooling limits of the package or system or cause excessive battery drain in portable devices.

In a prior Rambus dynamic random access memory (DRAM) system, illustrated in Fig. 1; three clock domains are used to control clock power use. Fig. 1 shows a memory core 11 and a memory interface 13. A first domain consists of the control

circuitry 15, and a second domain is the write path 17 into the memory core. A third path is read data path 19. Table 1 indicates which clock domains are on for different power modes. The clock power to all three domains can be turned off in a standby mode. The control domain is turned on to enter an active mode. The write data path 5 17 is additionally turned on when needed for a write operation, an active write mode. Read data path 19 is turned on for a read operation, an active read mode. In a read operation, the control domain is turned on first, to allow the control signals for the read to be provided to the memory core. Since there will be some delay before the data is available, the output data path for the data to be read need not have its clock 10 turned on until some period of time later. Thus, there is a latency between when the control logic is turned on and provided the clock signal to when the read data output path is turned on.

TABLE 1

power modes	Clock Domains		
	rclk	sclk	tclk
standby	off	off	off
active	on	off	off
active read	on	off	on
active write	on	on	off

In a prior Rambus system, the read data path is turned on automatically by the control logic when the control logic is activated for a read. Thus, a separate control signal does not need to be sent over the interface to turn on the read data path. A 25 register will store a count corresponding to the latency from a RAS control signal to when read data will be available from the memory core, and this register value is used to control the latency of the clock turn-on for the read data path.

One disadvantage of the prior Rambus system is the additional latency required for turning on the control logic to exit the standby power mode. Since the 30 interface control logic and datapath must be on before an incoming command can be processed and a memory operation started, the turn-on latency of the control logic and datapath directly adds to the memory access latency. This provides a power versus latency trade off.

Another method of limiting clock power consumption is to use a slower clock signal. This is done in microprocessors which have a low power or sleep mode. Typically, these are used in laptop computers in which the user can select a lower power mode, or the laptop will automatically enter the lower power or sleep mode in
5 the absence of any user input within a predefined period of time.

SUMMARY OF THE INVENTION

The present invention provides a memory device with multiple clock domains.

- 10 Separate clocks to different portions of the control circuitry create different clock domains. The different domains are sequentially turned on as needed to limit the power consumed. The turn on time of the domains is overlapped with the latency for the memory access to make the power control transparent to the user access of the memory core.
15 In one embodiment, the present invention separates out the RAS control logic into a separate clock domain from the CAS control logic. This smaller amount of RAS control logic can then be left on in a standby power mode to eliminate any visible latency from a RAS signal through to data access.

20 The write and read data paths are also in separate clock domains to further conserve power depending upon whether an operation is a read or write operation.

In one embodiment, the power control is implicit and transparent to the user. In a standby mode, a RAS signal will cause the control logic associated with the RAS control logic to activate the CAS clock domain an appropriate latency after receipt of the RAS signal without any further control over the memory interface required. When
25 a CAS read or write signal is received, that will implicitly, or automatically, cause the read or write clock domain to be turned on an appropriate latency after the CAS signal.

In yet another embodiment of the invention, the memory device can dynamically switch between a fast and a slow clock depending upon the needed data bandwidth. The data bandwidth across the memory interface can be monitored by the memory controller, and when it drops below a certain threshold, a slower clock can be used. The clock speed can be dynamically increased as the bandwidth demand increases. Thus, rather than a coarse switch between a slow or fast clock

speed depending upon user activity, clock speed can be switched automatically depending upon data access bandwidth requirements.

For a further understanding of the nature and advantages of the invention, reference should be made to the following description taken in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram of a prior art memory system with multiple clock domains.

Fig. 2A is a block diagram of a memory system according to one embodiment of the invention.

Fig. 2B is a diagram illustrating memory operation pipelining in accordance with the invention.

Fig. 3 is a block diagram of one embodiment of the segmentation of the control logic into clock domains according to the access pipelining of Fig. 2.

Fig. 4 is a block diagram of an alternate embodiment of Fig. 3 wherein the control logic is broken into two clock domains.

Fig. 5 is a block diagram of multiple control and data access clock domains according to one embodiment of the invention.

Fig. 6 is a timing diagram illustrating the sequencing of the clock domains of Fig. 5.

Figs. 7, 8 and 9 illustrate three embodiments for dynamically varying the clock speed in accordance with data bandwidth.

25

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Fig. 2A shows a memory system 10 which is controlled by a controller 12 over an interconnect bus 14 to a series of RDRAMs (Rambus dynamic random access memories) 16. The interconnect bus 14 includes two nine bit data busses 18 (BusDataA and BusDataB), and an eight bit primary control bus (RQ) 20. A clock 22 is provided in one direction along a transmit line 24 (TClk), and loops back along a receive clock line 26 (RClk).

In addition a low power, serial daisy-chained control interface is provided with daisy-chained segments 28 and a return segment 30. In other embodiments, the control sideband 28, 30 may be a bus instead of daisy-chained. In a powered down or nap mode, the primary control bus and data busses can be turned off with communication being initiated using the control sideband 28, 30.

5 Each RDRAM 16 includes multiple banks 32 of memory, each with associated core control logic 34. In addition, each chip includes interconnect logic 36.

Fig. 2B illustrates the pipelining of memory operations used in the present invention. As can be seen, the memory operations are broken up into five steps: 10 precharge, sense, transfer, data and close. By breaking up the control and data transfer of a memory access operation in this manner, the operations can be pipelined to increase throughput. For a further description of the pipelining according to the invention, reference should be made to co-pending application Serial No. 09/xxx, entitled "Pipelined Memory Device", filed even date herewith and assigned to 15 the same assignee, incorporated herein by reference.

Fig. 3 is a diagram illustrating one embodiment of interconnect logic 36 in a RDRAM. As can be seen, the control logic is broken up into segments, corresponding to the pipeline stages of Fig. 2B. In particular, there is a pre-charge input 38, a sense input 40, a transfer input 42, and a close input 44. The data 20 pipeline of Fig. 2B is not shown, since Fig. 3 shows only control logic. Fig. 3 additionally shows a retire input 46. This retire input is used for a two-stage write operation, in which data is first written to an internal buffer, and is then written from the buffer into the memory core upon receipt of a "retire" command from the memory controller.

25 Each of the control inputs is provided on a serial line which is converted into parallel with a respective converter 48. The control signals are sent as a packet, which is decoded in respective packet decode logic 50. From there, decoded control signals are sent to either a row state machine 52, or a column state machine 54. The state machine will then send the appropriate control signals to the memory core 30 control logic.

Not all of the control logic segments shown in Fig. 3 need to be operating at the same time. Even in a tightly pipelined operation, if reads and writes are alternating, for example, only one of the column read and column write control logic needs to be operating each cycle. Accordingly, each of the segments of Fig. 3 can be

made a separate clock domain, with the clock to the domain turned on and off as necessary.

The turning off of clock domains as they are not needed significantly reduces power consumption and heat generation of the memory chip. As will be described below with respect to Figs. 5 and 6, turning on and off of each clock domain is done with timing such that it is transparent to a memory access. In other words, the memory access will take no more time to complete than if all the clocks were left on all the time. An example of how the latency of the clock domain turn on is hidden behind the normal latency of the pipeline is set forth below with respect to Figs. 5 and 6.

In one embodiment, a sense operation turns on the clock domain for both the precharge control and close logic. A transfer operation (by the column control logic, also sometimes called the CAS control logic) turns on the retire logic (i.e., labeled the column write control logic in Fig. 3). A precharge operation can turn on the column control logic, or parts of it. Signal line 55 in Figs. 3 and 4 indicates that the control column logic activates the precharge control logic so as to restart precharging (i.e., of the sense amplifiers in the memory core) after the column state machine 54 has completed transfer of information between the sense amplifiers and the data input/output buffers.

Implicit control is also used to turn off clock domains not needed for later steps in a particular operation sequence. For example, a close operation can turn off secondary control domains, such as the transfer and retire logic in the column state machine.

Fig. 4 shows an alternate embodiment of the invention in which the control logic 36 is broken down into two clock domains, sense clock domain 56 and column clock domain 58. Row clock domain 56 receives and processes row control packets while column clock domain 58 receives and processes column control packets. Row control packets specify a row address and bank address for a sense operation (also known as a RAS operation), while the column control packets indicate column address and the type of column access operation (e.g., read or write), also known as a CAS operation, to be performed. Such an embodiment reflects a balance between the complexity and cost of more clock domains and the benefits of additional power savings. This embodiment recognizes that the major power savings is achieved by separating the RAS and CAS control operations.

Fig. 5 is a block diagram of one embodiment of the invention illustrating the different clock domains. Fig. 5 shows a RAS clock domain 60 including control logic 62 for the RAS signal and PCH(precharge) control logic. In addition, a clock receiver 64 is shown which connects to control bus line 66.

5 A second, CAS clock domain 68 includes CAS control logic 70 and a receiver 72. A write data clock domain 74 includes a write input pipeline 76 and a receiver 78. A read clock domain 80 includes a read output pipeline 82.

10 In the middle of Fig. 5 is shown a clock receiver 84 for the main clock used for the main control bus, lines 24 and 26 of Fig. 2A. This clock is synchronized with a delay lock loop (DLL) 86. Finally, a sideband control logic block 88 is shown.

15 In operation, sideband control logic 88 will always be on, even in a power down mode. In a power down mode, DLL 86 and clock receiver 84 can be turned off via a control line 90. Because of the significant latency required to turn on clock receiver 84 and DLL 86, an intermediate power mode, called a "nap" mode is used in which receiver 84 is on, but DLL 86 is in a low power configuration that provides faster synchronization time than when the DLL is entirely off. When in the nap mode, a memory access can be initiated with an appropriate control signal over sideband control line 28.

20 In an alternate embodiment, the sideband control logic 88 also receives RAS, or RAS and CAS control data for initiating the memory access, since clock domain 60 will be off when the memory access is initiated. When the sideband control logic 88 receives RAS, but not CAS, control data, the sideband control logic 88 will turn on CAS clock domain 68 so that it is ready when the subsequent CAS signal is sent.

25 A standby power mode is initiated by an appropriate control packet to sideband logic 88. This will cause DLL 86 to transition from the nap mode to an ON mode via a control signal on line 92. At the same time, RAS clock domain 60 will be turned on by the same line 92 by enabling the gating of the clock signal through AND gate 94. Thus, in standby mode, clock receiver 84 and DLL 86 are on as well as RAS clock domain 60. When a RAS signal is received, it can be immediately processed by RAS logic 62. In anticipation of the subsequent CAS signal, control logic 62 will provide a signal on a line 96 to turn on CAS clock domain 68. This will automatically, or implicitly, turn on the CAS clock domain. Thus, the user need not send special commands to turn on and off the clock domains but rather it is done automatically in accordance with the initiation of a memory access at the appropriate time. The signal

on line 96 is sent after a period of time corresponding to the latency of the CAS signal with respect to the RAS signal. This period of time can be stored in a register in RAS control logic 62 or can be determined by the memory controller. The signal on line 96 activates receiver 72, and also turns on the CAS clock domain 68 by activating AND gate 98.

5

The registers for setting the latency for implicit control of clock power domains can be programmed upon device initialization, or can be dynamically varied.

Alternately, instead of a register, the latency can simply be designed into the circuit.

When a CAS signal is received, it will either be a CASR (CAS read) or a
10 CASW (CAS write) signal. If it is a CASW, CAS control logic 70 will activate the write data domain 74 via control line 100, after an appropriate latency corresponding to when the write data will be received from bus 18. Control line 100 turns on receiver 78, and also enables AND gate 102 to provide the clock signal to write data domain 74.

15

Similarly, for a CASR operation, CAS control logic 70 activates AND gate 106, via a signal on line 104, to provide a clock signal (TCLK) to read clock domain 80.

20

Thus, in the present invention, when the main control busses are in standby mode, memory accesses can be performed over the fast main control bus, with power mode transitions being implicitly controlled to make the power mode transitions transparent. In this way, core access latencies are overlapped with power turn-on latencies. By automatically turning on certain interface power modes when a particular core operation command is given, implicit power control is provided which minimizes latency and saves valuable control bandwidth (no explicit power control is needed). This will efficiently integrate power control into the memory access pipeline
25 stream.

stream.

Fig. 6 illustrates these overlapped latencies. As shown in Fig. 6, the system of Fig. 5 is initially in a standby mode. When a RAS signal 110 is received, it will initiate the core row access as indicated by arrow 112. At the same time it will turn on the CAS clock domain as indicated by arrow 114. After a turn-on time and programmed latency, the CAS clock will be on at a point in time 116. At this point, the device is in what is considered an "active" mode. This turn-on time 116 is in advance of receiving a CASR signal 118 in its normal, pipeline timing. CAS-R signal 118 initiates a column read as indicated by arrow 120. The column read is initiated at a time tRCD after the row access, which is the latency of the row access. As also shown by an arrow 122,

30

CASR signal 118 will also turn on the read data path after a time delay tDAC corresponding to the latency from the column read until the data is available. Thus, at a point in time 124, the read data path will be on (clock domain 80 of Fig. 5). The device will now be in “active read” mode. Again, the latency from the CAS-R signal to 5 the read data path turn on time 124 matches the latency between the column read and when data is available as indicated by arrow 126.

For a write operation, similar latencies apply, as illustrated by CAS-W signal 128. This will initiate a column write operation as illustrated by arrow 130, and concurrently will provide a control signal as illustrated by line 132 to turn on the write 10 data path. For a write operation, the data is coming in, and thus, the write data path must be turned on at a time 134 in advance of when the write data 136 is received over the data bus. Again, the latency of the power domain turn on of the write data path is transparent or hidden behind the latency of the write data arriving. The write data arriving is also pipelined so that it is not provided until needed for a column write 15 in the memory core, as illustrated by arrow 137.

Slow Clock

Fig. 7-9 illustrate embodiments of the invention which allow the clock provided over the main control bus to be dynamically varied from fast to slow in accordance 20 with the data bandwidth requirements. In addition, a slow clock could be used to initiate a memory access operation without requiring the DLL 86 of Fig. 5 to be on, since the slower clock may not need synchronization. The slow clock would enable the access to be started concurrently with the CAS control clock domain being turned on.

25 A slower clock speed results in lower power because the AC power to switch the capacitive load connected to the clocks is reduced proportionately to the frequency. Also, at reduced clock speeds, the device may not require a high power phase compensation circuit (DLL or PLL), which is typically required for high speed operation. Depending on the slow clock frequency, the interface may operate without 30 phase compensation or use a lesser performance phase compensation circuit which consumes less power.

Fig. 7 illustrates one embodiment in which a memory 140 has an interface 142 with a slow clock input circuit 144 and a fast clock input circuit 146. Each input circuit

includes a clock receiver 148 and a DLL or PLL 150. A multiplexer 152 selects whether a fast or slow clock is being used via a select control line 154.

The clock source 155 itself provides both a fast clock and slow clock through a multiplexer 156.

- 5 A circuit 158 for selecting between the slow and fast clocks is preferably provided either in the controller or in a memory interconnect or some other device connected to the bus.

By monitoring bus traffic, the amount of bandwidth being used is determined by a monitor circuit 158. When bus traffic exceeds a predefined threshold (i.e., more 10 than a predefined number of memory accesses per second), the monitor circuit selects the fast clock, and otherwise it selects the slow clock. Depending on which clock is used, the fast or slow DLL or PLL and the unused receiver are preferably turned off to conserve power. In one embodiment, the monitor circuit 158 may be implemented using a programmed microprocessor, external to the memory circuit.

- 15 As referred to above, on a transition from a slow clock to a fast clock usage, during the latency of the fast clock turn on, operations can be occurring using the slow clock.

Fig. 8 is a variation of the circuit of Fig. 7 where two separate clock inputs are used, rather than a single, multiplexed line for the clock input.

- 20 Fig. 9 illustrates an alternate embodiment in which a single external fast clock is provided, which is internally divided using a clock divider circuit 160 to produce the slow clock.

As will be understood by those skilled in the art, the present invention may be embodied in other specific forms without departing from the spirit or central 25 characteristics thereof. Accordingly, the foregoing description is intended to be illustrative, but not limiting, of the scope of the invention which is set forth in the following claims.